REGULATORY DOCUMENT

# Accountability & Governance Framework Overview

Roles, Escalation Logic, and High-Level Schemas for Regulators

| VERSION | EFFECTIVE DATE | CLASSIFICATION | ISSUING ENTITY |
| --- | --- | --- | --- |
| 1.0 | December 29, 2025 | For Regulatory Review | TeraSystemsAI |

CONTENTS

## 1. Scope & Intent

### What This Framework Governs

- AI-assisted decision support systems deployed in regulated environments
- Assignment and traceability of responsibility for AI-influenced outcomes
- Escalation conditions requiring human review
- Audit and compliance artifact generation

### What This Framework Does Not Govern

- Model architecture, training methodologies, or algorithmic internals
- Data privacy or security controls (governed by separate policies)
- Procurement, vendor selection, or commercial terms

> AI systems may assist decisions. AI systems may never hold final authority. Responsibility is explicitly assigned and auditable at every decision point.

## 2. Foundational Governance Principle

> AI systems are decision-support tools. Legal, ethical, and operational responsibility always resides with humans or institutions.

This principle is:

- **Domain-agnostic:** Applies uniformly across healthcare, finance, legal, education, and enterprise contexts
- **Technology-independent:** Applies regardless of model type, vendor, or implementation approach
- **Structurally enforced:** Implemented through system architecture, not policy statements alone

The framework ensures this principle through deterministic role separation, mandatory escalation triggers, and immutable audit logging.

## 3. Explicit Role Model

The framework defines four immutable roles. These roles cannot be merged, bypassed, or reassigned during operation.

### AI System Role

- Pattern detection and analysis
- Risk estimation with uncertainty bounds
- Uncertainty quantification
- Recommendation generation

**Explicit Prohibitions:**

No final decisions. No authority. No enforcement. No responsibility acceptance.

### Human Reviewer Role

- Final decision authority
- Ethical and contextual judgment
- Right to override AI recommendations
- Accountability for approved outcomes

**Mandatory Requirement:**

This role cannot be bypassed. Required for all high-risk outcomes.

### Policy Configuration Role

- Defines confidence thresholds
- Defines escalation conditions
- Defines compliance constraints
- Defines domain-specific rules

**Governance Requirement:**

All rules are deterministic, versioned, and auditable.

### Institution Role

- Legal responsibility ownership
- Liability acceptance
- Governance enforcement
- Regulatory compliance

**Non-Transferable:**

Institutions cannot transfer accountability to AI systems.

## 4. Escalation Logic

### Escalation Triggers

Escalation to human review is mandatory when any of the following conditions are met:

| TRIGGER CATEGORY | CONDITION | RESPONSE |
|---|---|---|
| Confidence | Model confidence below policy threshold | Mandatory human review |
| Uncertainty | Uncertainty bounds exceed acceptable limits | Mandatory human review |
| Fairness | Bias or fairness constraints triggered | Mandatory human review |
| Domain Risk | Domain-specific risk conditions met | Mandatory human review |
| Override | Human requests review regardless of metrics | Review initiated |

### Escalation Sequence

1. **Trigger Detection:** System identifies escalation condition from policy rules
2. **Output Freeze:** AI recommendation is frozen; no further processing occurs
3. **Review Assignment:** Human reviewer is assigned with full context and uncertainty data
4. **Decision Logging:** Reviewer decision (approve, reject, modify) is immutably logged
5. **Responsibility Assignment:** Accountability record links outcome to human reviewer

> Escalation is deterministic, non-optional, and auditable. No configuration permits bypassing escalation when trigger conditions are met.

## 5. Responsibility Matrix

The Responsibility Matrix is a formal governance artifact that records accountability for every AI-assisted outcome. Each record contains:

| FIELD | DESCRIPTION | PURPOSE |
|---|---|---|
| Recommendation Source | AI system identifier and version | Traceability |
| Reviewer Identity | Human reviewer who evaluated output | Accountability |
| Approval Authority | Human or institution that approved outcome | Responsibility |
| Policy Version | Active policy rules at time of decision | Compliance |
| Escalation Record | Whether escalation occurred and why | Audit |
| Timestamp | Immutable timestamp of all actions | Verification |

### Matrix Properties

- **Immutable:** Records cannot be modified after creation
- **Logged:** All entries are written to append-only audit storage
- **Exportable:** Standard formats available for regulatory review (JSON, CSV, PDF)
- **Regulator-Ready:** Designed for external audit without requiring system access

## 6. Language & Claims Control

Language is treated as a governance surface. Phrases that imply AI authority create legal and compliance risk. The framework enforces language standards:

**PROHIBITED PHRASES**

- "The AI decided..."
- "The system approved..."
- "The algorithm determined..."
- "Automated decision..."

**REQUIRED ALTERNATIVES**

- "The AI recommended..."
- "The reviewer approved..."
- "The system identified..."
- "AI-assisted recommendation..."

### Rationale

- **Risk Reduction:** Eliminates language that could imply transferred responsibility
- **Legal Defensibility:** Maintains clear human accountability in all documentation
- **Audit Clarity:** Ensures consistent terminology across compliance artifacts

## 7. Cross-Domain Applicability

The governance structure applies uniformly across regulated domains. The foundational invariant does not change; only policies and thresholds vary.

**Healthcare**
Clinical decision support, diagnostics, treatment recommendations

**Finance**
Credit decisions, risk assessment, fraud detection

**Legal**
Case analysis, document review, risk scoring

**Education**
Assessment support, learning recommendations, admissions

**Enterprise**
HR decisions, procurement analysis, operational recommendations

**Government**
Benefits determination, regulatory analysis, public safety

> The accountability invariant is domain-agnostic. In every domain: AI recommends, humans decide, institutions are accountable.

## 8. What This Enables for Regulators

- ✓ Clear accountability chains for every AI-assisted outcome
- ✓ Deterministic escalation with auditable triggers
- ✓ Reduced automation bias through mandatory review
- ✓ Exportable compliance artifacts in standard formats
- ✓ Immutable audit logs for post-hoc review
- ✓ Language standards that preserve legal clarity
- ✓ Role separation that prevents responsibility diffusion
- ✓ Policy versioning for regulatory timeline reconstruction

### Compliance Verification

Regulators can verify framework compliance through:

- Responsibility Matrix export and audit
- Escalation log review
- Policy version history inspection
- Language compliance sampling

No access to system internals, source code, or model weights is required for compliance verification.

---